Chapter 4

System Architecture Overview

The routing node has two main architectural components:

Routing Engine—This component provides Layer 3 routing services and network management.

Packet Forwarding Engines—These high-performance, ASIC-based components provide Layer 2 and Layer 3 packet switching, route lookups, and packet forwarding.

The Routing Engine and the Packet Forwarding Engines perform their primary tasks independently, although they constantly communicate through a 100-Mpbs internal link. This arrangement streamlines forwarding and routing control and runs Internet-scale backbone networks at high speeds. Figure 19 shows the relationship between the Routing Engine and the Packet Forwarding Engines.

Figure 19: Routing Node Architecture

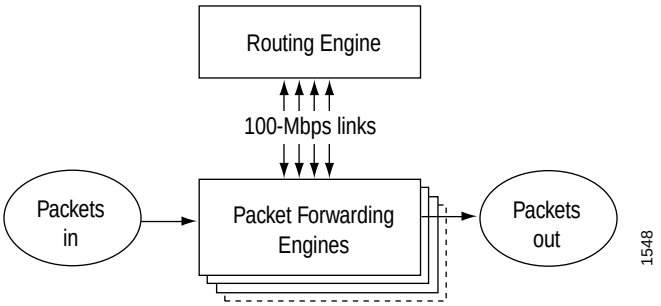

Routing Engine

The Routing Engine consists of JUNOS Internet software running on an Intel-based PCI platform. The JUNOS software is developed and optimized to scale to large numbers of network interfaces and routes. The software consists of a series of system processes running in protected memory modules on top of an independent operating system. The JUNOS kernel supports JUNOS system processes, which handle system management processes, routing protocols, and control functions (see Figure 20).

The Routing Engine has a dedicated 100-Mbps internal connection to the Packet Forwarding Engines.

Figure 20: Routing Engine Architecture

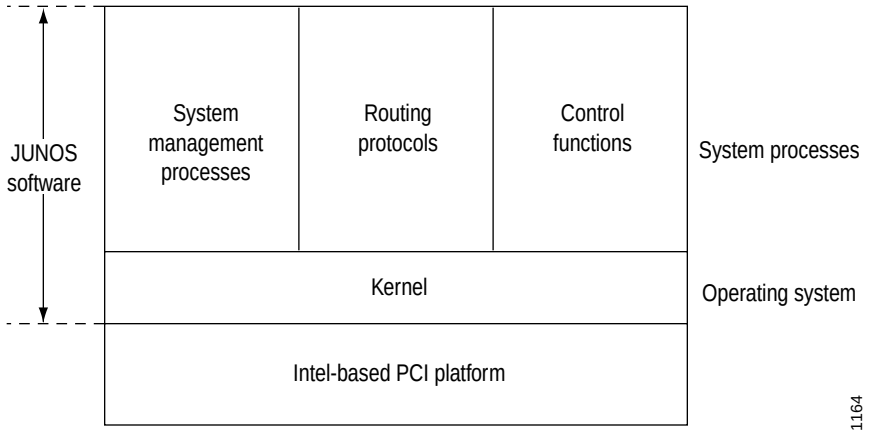

Routing Engine Functions

The Routing Engine handles all the routing protocol processes, as well as other software processes that control the routing node interfaces, system management, and user access to the routing node. These routing and software processes run on top of a kernel that interacts with the T-series Internet Processor in the Packet Forwarding Engine.

The Routing Engine provides the following functions:

Routing protocol packet processing—All routing protocol packets from the network are directed to the Routing Engine, and hence do not delay the Packet Forwarding Engine unnecessarily.

Software modularity—By dividing the different software functions into separate processes, the failure of one process is isolated from others and has little or no effect on them.

In-depth Internet functionality—Each routing protocol is implemented with a complete set of Internet features and provides full flexibility for advertising, filtering, and modifying routes. Routing policies are set according to route parameters (such as prefix, prefix lengths, and BGP attributes).

Scalability—The JUNOS routing tables are designed to hold all the routes in current and imminent networks. Additionally, the JUNOS software efficiently supports large numbers of interfaces and virtual circuits.

Management interface—Different levels of system management practices are provided, including a command-line interface (CLI) and SNMP.

Storage and change management—Configuration files, system images, and microcode can be held and maintained in primary and secondary storage systems, permitting local or remote upgrades.

Efficiency and flexibility monitoring—The routing node permits alarm handling and packet counting. For example, the routing node allows information to be gathered on every port, without adversely affecting packet forwarding performance.

The Routing Engine constructs and maintains one or more routing tables (see Figure 21). From the routing tables, the Routing Engine derives a table of active routes, called the *forwarding table*, which is copied into the Packet Forwarding Engines. The design of the T-series Internet Processor allows the forwarding table in the Packet Forwarding Engines to be updated without interrupting the routing node's forwarding.

Figure 21: Control Packet Handling for Routing and Forwarding Table Update

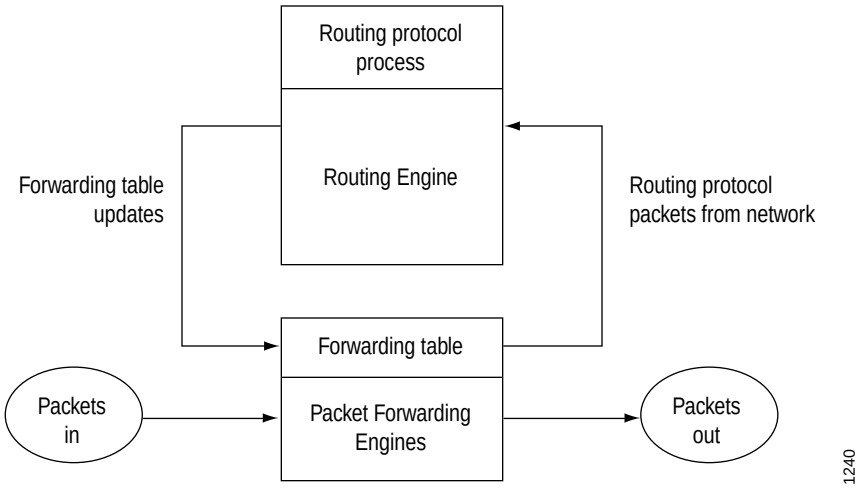

Packet Forwarding Engines

The Packet Forwarding Engines provide the Layer 2 and Layer 3 packet switching, forwarding, and route lookup functions. In a maximum configuration with eight FPC3s installed, the Packet Forwarding Engines can forward up to 640 million packets per second (Mpps) for all packet sizes. The maximum aggregate throughput rate for the routing node is 320 Gbps (full duplex). The Packet Forwarding Engines are implemented in ASICs that are physically located on the FPCs and the PICs.

Each Packet Forwarding Engine consists of the following components (see Figure 22):

Layer 2/Layer 3 Packet Processing ASIC, which performs Layer 2 and Layer 3 encapsulation and decapsulation, and manages the division and reassembly of packets within the routing node.

Queuing and Memory Interface ASICs, which manage the buffering of data cells in memory and the queueing of notifications.

T-series Internet Processor, which provides the route lookup function.

Switch Interface ASICs, which extract the route lookup key and manage the flow of data cells across the switch fabric.

Media-specific ASICs on the PICs that perform control functions tailored to the PIC media types.

Data Flow through the Routing Node

To ensure the efficient movement of data through the routing node, the routing node is designed so that ASICs on the hardware components handle the forwarding of data. Data flows through the routing node in the following sequence (see Figure 22):

Figure 22: Data Flow through the Routing Node

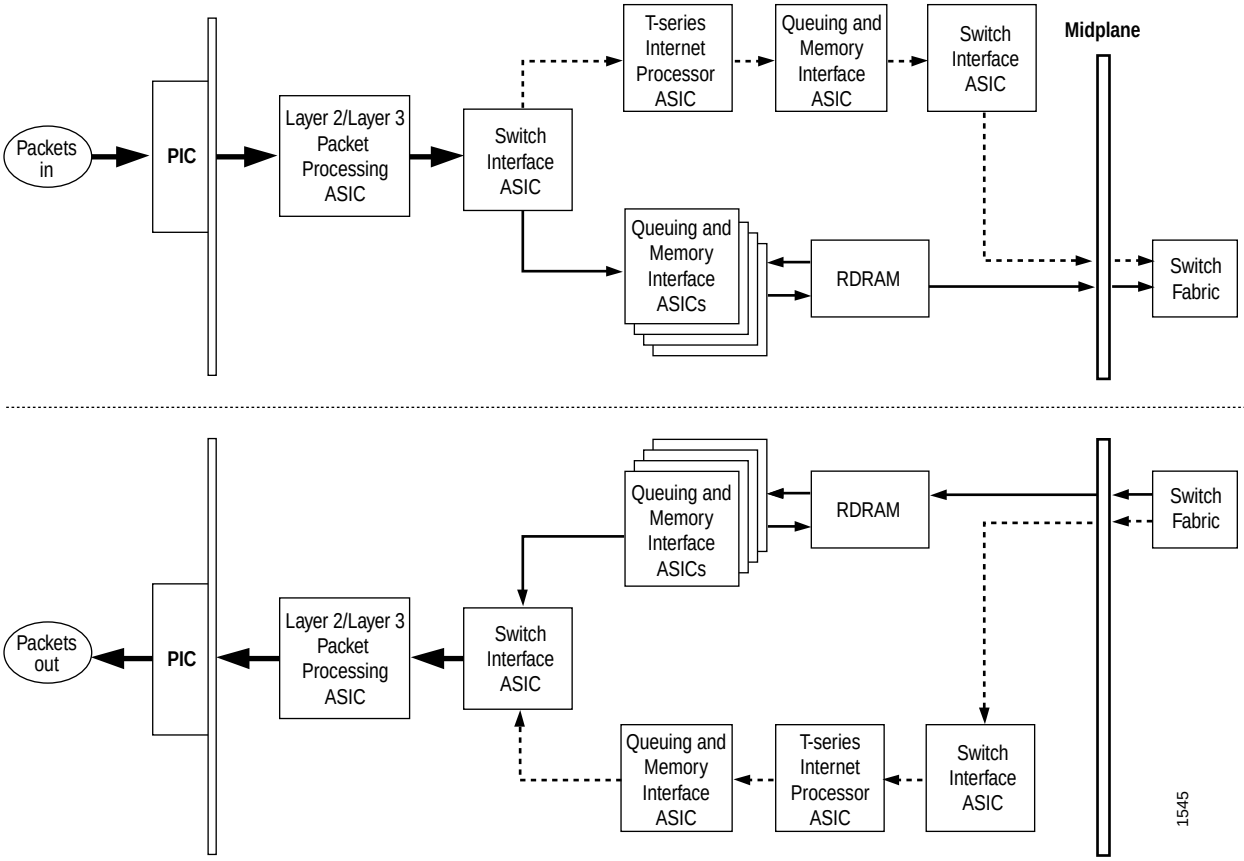

1. Packets arrive at an incoming PIC interface.
2. The PIC passes the packets to the FPC, where the Layer 2/Layer 3 Packet Processing ASIC performs Layer 2 and Layer 3 parsing and divides the packets into 64-byte cells.
3. The Switch Interface ASIC extracts the route lookup key, places it in a notification and passes the notification to the T-series Internet Processor. The Switch Interface ASIC also passes the data cells to the Queuing and Memory Interface ASICs for buffering.
4. The Queuing and Memory Interface ASICs pass the data cells to memory for buffering.
5. The T-series Internet Processor performs the route lookup and forwards the notification to the Queuing and Memory Interface ASIC.

6. The Queuing and Memory Interface ASIC sends the notification to the Switch Interface ASIC facing the switch fabric, unless the destination is on the same Packet Forwarding Engine. In this case, the notification is sent back to the Switch Interface ASIC facing the outgoing ports, and the packets are sent to the outgoing port without passing through the switch fabric (see Step 13).

7. The Switch Interface ASIC sends bandwidth requests through the switch fabric to the destination port. The Switch Interface ASIC also issues read requests to the Queuing and Memory Interface ASIC to begin reading data cells out of memory.

8. The destination Switch Interface ASIC sends bandwidth grants through the switch fabric to the originating Switch Interface ASIC.

9. Upon receipt of each bandwidth grant, the originating Switch Interface ASIC sends a cell through the switch fabric to the destination Packet Forwarding Engine.

10. The destination Switch Interface ASIC receives cells from the switch fabric. It extracts the route lookup key from each cell, places it in a notification, and forwards the notification to the T-series Internet Processor.

11. The T-series Internet Processor performs the route lookup, and forwards the notification to the Queuing and Memory Interface ASIC.

12. The Queuing and Memory Interface ASIC forwards the notification, including next-hop information, to the Switch Interface ASIC.

13. The Switch Interface ASIC sends read requests to the Queuing and Memory Interface ASIC to read the data cells out of memory, and passes the cells to the Layer 2/Layer 3 Packet Processing ASIC.

14. The Layer 2/Layer 3 Packet Processing ASIC reassembles the data cells into packets, adds Layer 2 encapsulation, and sends the packets to the outgoing PIC interface.

15. The outgoing PIC sends the packets out into the network.